

UNITED STATES UTILITY PATENT APPLICATION

FOR

A DYNAMIC NETWORK INTERFACE

Inventor(s):

Solomon Trainin

Prepared by:

BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN LLP

12400 Wilshire Boulevard, Seventh Floor
Los Angeles, California 90025-1026
(408) 720-8598

Attorney Docket No.: 42390.P10772

EXPRESS MAIL CERTIFICATE OF MAILING

"Express Mail" mailing label number EL617 209298 US
Date of Deposit June 27, 2009

I hereby certify that this paper or fee is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 CFR 1.10 on the date indicated above and is addressed to the Commissioner of Patents and Trademarks, Washington, D.C. 20231.

Kelli A. Ivey

(Typed or printed name of person mailing paper or fee)

Kelli Ivey

(Signature of person mailing paper or fee)

A DYNAMIC NETWORK INTERFACE

FIELD OF THE INVENTION

[0001] The invention described herein relates to the field of computer networks. More particularly, the invention relates to a network interface protocol.

BACKGROUND

[0002] High-speed computer network topologies, such as Local Area Networks (LAN) and Wide Area Network (WAN), demand increasing bandwidth and data throughput. Consequently, data traveling across high-speed computer networks may “bottleneck” once the data is received by a Network Interface Card (NIC) within a target computer system. In order to accommodate continued increases in LAN/WAN network bandwidth, bottlenecks within the network must be prevented.

[0003] One reason data may bottleneck in a NIC is because the network interfaces used in current NICs may be inefficient in the way in which they handle incoming data. For example, a prior art network interface handles incoming data by relying on a NIC software driver to retrieve data from a temporary receiving buffer and store the data into a protocol memory before storing the data to application memory. Network software applications may then access the data from the application memory. Such a network interface involves numerous intermediate copying operations that take time and may cause network throughput to bottleneck.

BRIEF DESCRIPTION OF THE DRAWINGS

[0004] The features and advantages will become apparent from the following detailed description in which:

[0005] Figure 1 illustrates a computer network according to one embodiment.

[0006] Figure 2 illustrates a Network Interface Card (NIC) according to one embodiment.

[0007] Figure 3 illustrates an RBD queue according to one embodiment.

[0008] Figure 4 illustrates the contents of an incoming frame of data.

[0009] Figure 5 illustrates fields that may be stored within a Receiving Frame Descriptor (RFD) corresponding to a received frame of data.

[0010] Figure 6 is a flow diagram illustrating the operation of a dynamic network interface according to one embodiment.

DETAILED DESCRIPTION

[0011] A dynamic network interface is described, intended to enable the efficient processing of received data within a computer network by a target computer system by reducing excessive copying of the received data prior to being accessed by a network software application.

A Computer Network

[0012] Figure 1 illustrates a computer network according to one embodiment. Computer systems may communicate over a network 105, such as a Local Area Network (LAN) or Wide Area Network (WAN), depending on the relative location and requirements of the computer systems within a network. Accordingly, data may be transmitted across a network via a wired connection, such as T1 or twisted pair, or by wireless means, such as a Radio Frequency (RF) signal.

[0013] A target system 115 is any computer system that is the destination of data transmitted across a network. A source computer system 110 may transmit data to the target computer system either after the data has been requested by the target computer system, or the source computer system may transmit data unsolicited by the target computer system. In either case, data sent to a target computer system by a source computer system is stored within the target computer system, at least temporarily, in order to allow software programs, such as a network software application, to access the data.

A Network Interface Sub-System

[0014] Data received by a target computer system may be copied from the target computer system to other locations within the receiving host computer system or

other computer system in order to make the received data accessible to data processing software or hardware, such as a network software application.

[0015] Data sent to a target computer system may be received by a network interface sub-system where it may be copied from a temporary storage device, such as buffer memory, to Application Memory (AM) in order to allow a network software application to access the data. AM is a location of memory that has been dedicated, either permanently or temporarily, to a network software application. Data stored within AM may then be accessed by a network software application running on the host computer system.

[0016] In one embodiment, the network interface sub-system is a Network Interface Card (NIC). Figure 2 illustrates an NIC according to one embodiment. In one embodiment, a NIC 205 consists of an embedded processor 225, embedded memory 230 and a receiving buffer 210. A receiving buffer may be used to temporarily store received data until the data can be copied or transferred to another location within the NIC or host computer system. Received data may then be copied to driver memory 220 where it may be accessed by software programs, such as a software driver, that control the NIC and subsequently copied to application memory 215. Application memory may be used to store the received data in order to allow a network software application to access it.

[0017] Device memory and application memory may exist within the same physical memory device or in separate memory devices. In one embodiment, application memory and device memory are allocated within system memory of the host target computer system, whereas the temporary receiving buffers are located in the NIC.

[0018] Once data is received by a NIC, an embedded processor may process the received data in order for a software application to access the data. For example, received data temporarily stored in a receiving buffer may be copied to device memory by a software driver running on the host system. The software driver may then copy the data to an application memory space where it may be accessed by a software application running on the host computer system.

A Dynamic Network Interface

[0019] A dynamic network interface is described herein, in which application software may access data received by a target computer system without waiting for the data to be copied by a software driver or other program to intermediate memory locations, such as device memory, before being stored to application memory.

[0020] In one embodiment, frames of data within a messaging flow, such as a Transmission Control Protocol (TCP) flow, are temporarily stored within a receiving buffer and then stored into buffers within application memory (AM) buffers that have been allocated by one or more network software applications. A messaging flow refers to one or more frames of data that are intended to be received by a particular computer system for use by a software application running on the receiving computer system.

[0021] A network software application may access received data frames stored in the AM buffers without having to wait for the data to be copied to an intermediate memory location and subsequently stored into application memory. Frames of data that are stored to application memory in such a manner are referred to as "Zero Copy" (ZC) frames, as the frames require no intermediate copying or

transferring prior to being stored within application memory. Accordingly, a portion of a messaging flow to which a ZC frame belongs is referred to as a ZC flow, while an AM buffer in which ZC frames are stored are referred to as a ZC AM buffer.

[0022] In one embodiment, frames received by the target computer system may be determined to be ZC frames by a device driver running on the host computer that controls the NIC hardware. Once the device driver has determined that the received frame may be processed as a ZC frame, embedded software running on the NIC may configure the frame as a ZC frame and copy the data from the receiving buffer to a ZC AM buffer or buffers.

[0023] After it is determined that a frame of data may be “zero copied” to application memory, one or more circular Receiving Buffer Descriptor (RBD) queues may be created and a predetermined number of RBD’s allocated within. Each RBD entry may contain at least a value corresponding to a start address and a value indicating the size of a ZC AM buffer to which an RBD corresponds. In one embodiment, the value corresponding to a start address is the start address of the corresponding ZC AM buffer within application memory. In other embodiments, the value may be an offset, sequence number, or other value from which the address of a ZC AM buffer may be derived. Furthermore, when a new ZC AM buffer is allocated within application memory, a new RBD entry may be stored within an available RBD queue entry or one that later becomes available.

[0024] Figure 3 illustrates an RBD queue according to one embodiment. A next available RBD queue entry may be indicated by a Head Pointer 305, whereas an oldest incomplete RBD queue entry may be indicated by a Tail Pointer 310. Once

a ZC AM buffer has been filled, an RBD corresponding to the filled ZC AM buffer may be indicated by a Last Completed RBD Number (LCRN) pointer 315. The LCRN pointer may be communicated to a NIC device driver in order to inform the device driver that the ZC AM buffer is full and will be de-allocated. The device driver may use this information to keep track of available AM buffers in order to process non-ZC frames. The RBD corresponding to the filled ZC AM buffer may then be de-allocated, allowing other pending or future RBD's to be stored within the queue. In one embodiment, a fulfilled RBD is de-allocated by invalidating the RBD through incrementing of the Tail Pointer.

[0025] Upon receiving a frame of data into the receiving buffers of a receiving host computer system's NIC, it may be determined whether the data is intended for the receiving host computer. In one embodiment, a network software application running on the target host computer system determines whether a received frame is intended for the target host computer and whether it is intended to be processed by the network software application. This information may be obtained by interpreting header information associated with the received frame.

[0026] Received frames may arrive in various network protocols, including TCP. In general, a flow of TCP frames may be identified by using various header parameters associated with the protocol. For example, in one embodiment, Internet Protocol (IP) Source Address, IP Destination Number, Destination Port Number, and Source Port number may be used to identify TCP frames of a particular flow as being intended for the receiving computer system.

[0027] When a new frame of data is received and the receiving host computer system determines that the frame is intended for the receiving host computer

system, it is then determined whether the frame may be treated as a ZC frame and therefore copied directly from receiving buffers within the receiving host computer NIC to application memory without intermediate copying steps. In one embodiment, a received frame of data within a TCP flow is “zero copied” to application memory if there is enough space for the frame within any of a group of allocated ZC AM buffers, and if not, new ZC AM buffers may be allocated.

[0028] If it is determined that the received frame may be “zero copied” to an existing ZC AM buffer within application memory then the data is copied directly into the available ZC AM buffer(s). If all allocated ZC AM buffers are unable to store the received frame and available space within the RBD queue exists, a new ZC AM buffer may be allocated and the corresponding RBD placed in the RBD queue, thereby enabling a new ZC flow of ZC frames.

[0029] In one embodiment, a flow descriptor is created that corresponds to a newly enabled ZC flow. A flow descriptor may contain flow control fields and RBD array status information pertaining to a ZC flow, which may be used by embedded software or other software/hardware within the target computer system to retrieve status information pertaining to the ZC flow or to identify the ZC flow.

[0030] Various fields may be stored within the flow descriptor using various bit lengths to represent each field. In one embodiment, the flow descriptor contains a 4-bit Flow Identification Number, a 1-bit Internet Protocol (IP) type, a 32 or 128-bit IP Source Address, a 32 or 128-bit IP Destination Address, a 16-bit Destination Port Number, a 16-bit Source Port Number, a 32-bit Start Sequence Number, a 1-bit TCP/UDP number, and an RBD Array Pointer. Any one or combination of these fields may be used to identify and recognize a particular

flow. In addition, the flow descriptor may contain an RBD Array indicating the number of RBD's within an RBD queue, and the number RBD's associated with frames within a ZC flow. Other fields may also be stored within the flow descriptor in other embodiments.

[0031] Once a ZC flow descriptor is created and at least one ZC AM buffer is allocated within application memory, ZC frames that were temporarily stored in receiving buffers may be stored in the allocated ZC AM buffers corresponding to a particular ZC flow. Once an ZC AM buffer is allocated, values corresponding to the buffer's start address and size are then stored within an available RBD, which is indicated by a Head Pointer within the RBD queue.

[0032] Figure 4 illustrates the contents of an incoming frame of data. Each received frame in a flow may contain a header 405 and a payload 410. The payload is the actual data that will be used by application software, whereas the header contains information relating to the particular received frame of data and the flow to which it belongs. After allocating an ZC AM buffer and storing the buffer's start address and size into an RBD, the payload of a received ZC frame is stored within the allocated ZC AM buffer. A header associated with the ZC frame is stored within a Receiving Frame Descriptor (RFD), which is stored in an RFD pool with other RFD's corresponding to other received frames. An RFD may be used to identify a particular frame.

[0033] Figure 5 illustrates fields that may be stored within an RFD corresponding to a received frame of data. Each newly received frame of data - ZC frame or otherwise - may store its header information in an RFD. In one embodiment, an RFD contains a 6-bit Last Complete RBD Number (LCRN) 505, a 4-bit Flow

Identification Number 510, a 32-bit Sequence Number 530, a 1-bit flag to indicate that a flow has started 515, a 1-bit flag to indicate that a flow has terminated 520, and a 2-bit number to indicate the reason for termination 525.

[0034] The 32-bit sequence number may indicate a start sequence number, if the start bit is set, and an end sequence number, if the terminated bit is set.

Furthermore, sequence numbers may identify relative position of a byte or group of bytes within a received frame of data in a particular flow. Accordingly, RFD structures may be used by a device driver controlling the NIC hardware to determine whether a frame of data has been configured as a ZC frame. A device driver may then determine whether the frame data should be ignored (if it is a ZC frame) or whether to retrieve the frame data (if the data is a non-ZC frame) and store it into device memory or some other memory from which it may be copied into application memory.

[0035] If a ZC bit 540 is set within an RFD, the device driver, in one embodiment, may also use the RFD to 1) check the flow corresponding to the flow ID within the RFD, 2) Release the RBD referenced by the LCRN pointer, 3) calculate the number of remaining valid RBD's in the RBD queue, and 4) determine whether new RBD content should be supplied to embedded software running on the NIC by determining whether the number of valid RBD's in the queue is less than some predetermined number.

One Embodiment of a Dynamic Network Interface

[0036] Figure 6 is a flow diagram illustrating the operation of a dynamic network interface according to one embodiment. After the host computer determines that the received data frame is intended for an application software program running

on the host computer 605, the header information associated with the received frame of data is stored within an RFD 610.

[0037] A determination may then be made as to whether the received frame may be “zero copied” directly to application memory from a receiving buffer. A received buffer may be “zero-copied” to application memory if available ZC AM buffer space exists 615. If not, 620 the received frame may be “zero-copied” to application memory if there is enough room in the RBD queue for a new ZC AM buffer to be allocated in application memory.

[0038] If the received frame cannot be stored into an existing ZC AM buffer and there is no space within the RBD queue for a new zero-copy AM buffer to be allocated within application memory, a prior art method 625 of handling the received frame is invoked. A prior art method may involve methods that do not support storing the received frame directly to application memory after it is received in a receiving buffer, but instead may invoke other intermediate copying operations. If there is enough space in an existing ZC AM buffer or buffers, the received frame data is stored in the available ZC AM buffer or buffers and the corresponding RFD zero-copy bit is set to indicate that the frame is a zero-copy frame 630.

[0039] In order to find an available ZC AM buffer to store the received data, a search is conducted among a pre-determined number RBD's, each corresponding to a particular ZC AM buffer. In one embodiment, the search is based on a sequence number and size of the received frame data stored within each RBD. The sequence number is used to calculate an AM buffer's offset address and the data size is used to determine whether there is enough space in the ZC AM buffer

for received data. A successful search identifies an ZC AM buffer address derived from the frame sequence number that is able to contain the received frame data. A search fails and is aborted if no such ZC AM buffer can be identified.

[0040] In one embodiment, the search may be aborted if the number of searched ZC AM buffers exceeds the predefined number, thus aborting the ZC flow itself. If the ZC flow is aborted, this may be reflected within the RFD associated with the received frame by setting a termination bit and an appropriate termination reason field.

[0041] If it is determined that there is no space for the received frame data in existing zero-copy ZC AM buffers, a new ZC AM buffer may be allocated to store the received frame data as a zero-copy frame. In one embodiment, a new ZC AM buffer is allocated if there is an available RBD queue entry in which to store an RBD associated with the newly allocated ZC AM buffer . If the RBD queue is full, prior art methods are used to copy the received data to application memory that may involve intermediate copying steps. If there is available space within the RBD queue, a new zero-copy AM buffer is allocated 635 and a corresponding RBD is placed within the RBD queue. The data may be stored within the newly-allocated ZC AM buffer and the RFD updated to reflect that the received frame is a ZC frame 640. A ZC flow descriptor may also be created to indicate that a new ZC flow has begun 645.

[0042] When a ZC AM buffer is fulfilled 650, a corresponding LCRN pointer within the RBD queue may be used to inform a software program, such as a NIC device driver, that the corresponding RBD queue entry is available to receive a new RBD. Accordingly, this information may be communicated to software

controlling the zero-copy operation, such as embedded software running on the NIC, to invalidate the RBD corresponding to the fulfilled AM buffer so that the RBD queue entry may be used to hold future or pending RBD's 655.

[0043] While this invention has been described with reference to illustrative embodiments, this description is not intended to be construed in a limiting sense. Various modifications of the illustrative embodiments, as well as other embodiments, which are apparent to persons skilled in the art to which the invention pertains are deemed to lie within the spirit and scope of the invention.

FOIA b 7 - D